

Certificate of E-Filing
I hereby certify that this document is being
transmitted to the Patent and Trademark Office
via EFS-Web on June 4, 2008.

/Joseph G. Swan/
Joseph G. Swan

PATENT
Atty. Docket No. 200210109-1

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Application of:

KEITH FARKAS, ET AL.

Serial No.: 10/621,067

Filed: July 16, 2003

For: Heterogeneous Processor Core
Systems for Improved Throughput

Group Art Unit: 2195

Examiner: Tang, Kenneth

Conf. No.: 1252

APPEAL BRIEF
ON APPEAL TO THE BOARD OF PATENT APPEALS AND INTERFERENCES

Mail Stop Appeal Brief - Patent
Commissioner for Patents
P.O. Box 1450
Alexandria, Virginia 22313-1450

Appellants in the above-captioned patent application appeal the final rejection of
claims 1, 2, 5-8, 11, 12, 15 and 17-24 set forth in the Office Action mailed January 17,
2008, a Notice of Appeal having been timely filed on April 16, 2008.

I. REAL PARTY IN INTEREST

The real party in interest in this application is Hewlett-Packard Development Company, L.P., pursuant to an assignment recorded on December 2, 2003, at reel 014168, frame 0694.

II. RELATED APPEALS AND INTERFERENCES

Appellants are not aware of any related appeals, interferences or judicial proceedings.

III. STATUS OF CLAIMS

Claims 1, 2, 5-8, 11, 12, 15 and 17-24 have been finally rejected and are the subject matter of this appeal. Claims 3, 4, 9, 10, 13, 14 and 16 have been canceled. In accordance with 37 C.F.R. § 1.192(c)(9), a copy of the claims involved in this appeal is included in the Claims Appendix attached hereto.

IV. STATUS OF THE AMENDMENTS

No amendment has been filed subsequent to the final rejection.

V. SUMMARY OF CLAIMED SUBJECT MATTER

The present invention concerns systems, methods and techniques for improving computer system processing performance. In the preferred embodiments, multiple different kinds of processor cores are provided and computer processing jobs are moved or transferred among them in order to improve a throughput metric.

In particular, independent claim 1 is directed to a computer system that includes multiple computer processor cores (e.g., as described at page 3 lines 20-22 and 35-36 of the Specification) in which at least two differ in processing performance (e.g., as described at page 4 lines 5-10 of the Specification), and in which all execute the same instruction set (e.g., as described at page 4 lines 1-5 of the Specification). The computer system also includes a performance measurement and transfer mechanism that moves a plurality of executing computer processing jobs amongst the computer processor cores to improve a throughput metric. See, e.g., the Title, page 1 lines 10-14, page 2 lines 10-11, page 11 lines 14-24 and page 12 lines 17-22 of the Specification.

Independent claim 7 is directed to a method for operating multiple processor cores, in which multiple computer processor cores are placed on a single semiconductor die (e.g., as described at page 3 lines 20-22 and 35-36 and page 15 lines 3-5 of the Specification), in which at least two computer processor cores differ in processing performance (e.g., as described at page 4 lines 5-10 of the Specification), and in which all execute the same instruction set (e.g., as described at page 4 lines 1-5 of the Specification). Performance is measured for each of a set of computer processing jobs hosted amongst the plurality of computer processor cores (e.g., as described at page 4 lines 31-35, page 5 lines 21-24 and page 8 line 29 through page 9 line 25 of the Specification), and individual ones of the computer processing jobs are transferred amongst targeted ones of the computer processor cores to improve a throughput metric (e.g., as described in the Title, page 1 lines 10-14, page 2 lines 10-11, page 11 lines 14-24 and page 12 lines 17-22 of the Specification).

VI. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

Claim 1 stands rejected on the ground of nonstatutory obviousness-type double patenting over claim 1 of U.S. Patent No. 7,093,147 (the '147 patent) in view of U.S. Patent Publication No. 2003/0110012 (Orenstien); claim 7 stands rejected on the ground of nonstatutory obviousness-type double patenting over claim 14 of the '147 patent in view of Orenstien; claims 1, 2, 5-8, 11, 12, 15, 17, 19, and 22 stand rejected under 35 U.S.C. § 102(e) over Orenstien; claims 18 and 23 stand rejected under 35 U.S.C. § 103(a) over Orenstien in view of an article titled, "Trends in Network and Pervasive Computing - ARCS 2002", April 2002 (Schmeck); claims 20-21 stand rejected under § 103(a) over Orenstien in view of U.S. Patent 5,913,068 (Matoba); and claim 24 stands rejected under § 103(a) over Orenstien in view of U.S. Patent 6,986,141 (Diepstraten).

VII. ARGUMENT

Authority Pertaining to Issues on Appeal

Anticipation Rejections Under 35 USC § 102

The requirements for showing anticipation under § 102 are described in M.P.E.P. § 2131 as follows:

"A claim is anticipated only if each and every element as set forth in the claim is found, either expressly or inherently described, in a single prior art reference." Verdegaal Bros. v. Union Oil Co. of California, 814 F.2d 628, 631, 2 USPQ2d 1051, 1053 (Fed. Cir. 1987).

With respect to a § 102 rejection, the Federal Circuit also has held that "The identical invention must be shown in as complete detail as is contained in the . . . claim." Richardson v. Suzuki Motor Co., 868 F.2d 1226, 1236, 9 U.S.P.Q.2d 1913, 1920, (Fed.

Cir. 1989). "For a prior art reference to anticipate in terms of 35 U.S.C. § 102, every element in the claimed invention must be shown in a single reference." In re Bond, 15 USPQ2d 1566, 1567 (Fed. Cir. 1990) (quoting Diversitech, 7 USPQ2d 1315). In addition, that single reference must arrange the elements exactly as in the claim under review, although identity of terminology is not required. *Id.*

In addition, when inherency is asserted extrinsic evidence must be cited to show that the missing descriptive matter is necessarily present in the thing described in the reference:

To establish inherency, the *extrinsic evidence* [emphasis added] "must make clear that the missing descriptive matter is necessarily present in the thing described in the reference, and that it would be so recognized by persons of ordinary skill." Continental Can Co. v. Monsanto Co., 948 F.2d 1264, 1268, 20 U.S.P.Q.2d 1746, 1749 (Fed. Cir. 1991). "Inherency, however, may not be established by probabilities or possibilities. The mere fact that a certain thing may result from a given set of circumstances is not sufficient." *Id.* at 1269, 20 U.S.P.Q.2d at 1749 (quoting In re Oelrich, 666 F.2d 578, 581, 212 U.S.P.Q. 323, 326 (C.C.P.A. 1981)).

In re Robertson, (Fed. Cir. 1999) 169 F.3d 743, 745; 49 U.S.P.Q.2d 1949.

Obviousness Rejections Under 35 USC § 103

The Supreme Court has set forth the following general standard with respect to any determination of obviousness:

"Under § 103, the scope and content of the prior art are to be determined; differences between the prior art and the claims at issue are to be ascertained; and the level of ordinary skill in the pertinent art resolved. Against this background the obviousness or nonobviousness of the subject matter is determined. Such secondary considerations as commercial success, long felt but unsolved needs, failure of others, etc., might be utilized to give light to the circumstances surrounding the origin of the subject matter sought to be patented."

Graham v. John Deere Co. of Kansas City, 383 U.S. 1, 17-18, 86 S. Ct. 684, 15 L. Ed. 2d 545 (1966), quoted approvingly by *KSR Int'l Co. v. Teleflex Inc.*, 127 S. Ct. 1727, 1734 (U.S. 2007).

When performing this analysis, all claim limitations must be considered. See, e.g., MPEP § 2143.01. At the same time, the analysis requires a determination as to whether the claimed invention “as a whole” would have been obvious just before the claimed invention was made to a person of ordinary skill in the art. See, e.g., MPEP § 2142.

It is noted that, “rejections on obviousness cannot be sustained with mere conclusory statements...” MPEP § 2142, quoting *In re Kahn*, 441 F.3d 977, 988, 78 USPQ2d 1329, 1336 (Fed. Cir. 2006), which in turn was quoted approvingly by the Supreme Court in *KSR International Co. v. Teleflex Inc.*, 82 USPQ2d 1385, 1395-97 (2007). In addition, “impermissible hindsight must be avoided and the legal conclusion [regarding obviousness] must be reached on the basis of the *facts* gleaned from the prior art [emphasis added].” MPEP § 2142.

More specifically, “the examiner must provide *evidence* which as a whole shows that the legal determination sought to be proved (i.e., the reference teachings establish a *prima facie* case of obviousness) is more probable than not [emphasis added].” MPEP § 2142.

Finally, even where all of a claim’s limitations can be found in the prior art, the examiner must provide a convincing reason as to why one of ordinary skill in the art would have been prompted to combine such limitations in the same manner as recited in claim.

"Although common sense directs one to look with care at a patent application that claims as innovation the combination of two known devices according to their established functions, it can be important to identify a reason that would have prompted a person of ordinary skill in the relevant field to combine the elements in the way the claimed new invention does. This is so because inventions in most, if not all, instances rely upon building blocks long since uncovered, and claimed discoveries almost of necessity will be combinations of what, in some sense, is already known."

KSR Int'l Co. v. Teleflex Inc., 127 S. Ct. 1727, 1741 (U.S. 2007).

Double-Patenting Rejection over claim 1 of the '147 patent in view of Orenstien

Claim 1

Independent claim 1 is directed to a computer system that includes multiple computer processor cores in which at least two differ in processing performance, and in which all execute the same instruction set. The computer system also includes a performance measurement and transfer mechanism that moves a plurality of executing computer processing jobs amongst the computer processor cores to improve a throughput metric.

The foregoing combination of features is not disclosed or suggested by the asserted combination. For instance, no permissible combination of claim 1 of the '147 patent and Orenstien would have disclosed or suggested at least the feature of moving a plurality of executing computer processing jobs amongst a plurality of computer processor cores to improve a throughput metric.

The Examiner acknowledges that this feature is neither disclosed nor suggested by claim 1 of the '147 patent, but then goes on to argue that Orenstien discloses such a feature. Specifically, the Examiner argues that,

"Orenstien teaches a computer system monitoring the operational activity of its dual-core processor 200 to provide load balancing between the cores by migrating processes from one core to the other based on a performance metric ([0017], [0022]-[0023]). Orenstien teaches that the performance metric, for example, can be based on power, or load for load balancing, or other things (page 2, last line of [0021]). The result of this migration based on the monitoring of the performance measurement is an increase in throughput or clock frequency for each core ([0020], [0026], [0037], [0039])." [Emphasis in original]

In this regard, Orenstien discusses techniques for modifying certain parameters pertaining to process execution based on power-consumption and thermal considerations. However, it does not say anything at all about the above-referenced feature of the present invention.

At the outset, it is noted that even the Examiner does not identify a single "throughput metric" that is even allegedly disclosed or suggested in Orenstien, and Appellants are unable to find any such disclosure in Orenstien. Accordingly, no permissible combination of claim 1 in the '147 patent and Orenstien possibly could have disclosed or suggested the presently recited feature of moving a plurality of executing computer processing jobs amongst a plurality of computer processor cores to improve any throughput metric.

In addition, all of the specific portions of Orenstien cited by the Examiner pertain to the modification of processing parameters (e.g., individual processor power, individual processor clock frequency and/or processing load distribution) based on *power consumption and/or temperature*. None of these specifically cited portions of Orenstien says or even suggests anything at all about any "throughput metric", as presently recited.

The relevant parts of those specifically cited portions of Orenstien are briefly summarized as follows: paragraph [0017] discusses redistributing processing tasks

based on power-consumption and/or thermal considerations; paragraphs [0022]-[0023] generally discuss different kinds of processors and processor cores; the last line of paragraph [0021] talks about using a "power consumption metric" in order to level the load between processors; paragraph [0020] discusses the use of a power-consumption or temperature metric for load-balancing purposes; paragraph [0026] mentions the use of a "power monitor 260"; paragraph [0037] discusses the testing of certain conditions pertaining to power consumption and then controlling voltage and/or frequency to the processors based on the results; and paragraph [0039] of Orenstien discusses:

"[shifting processing] between units, cores, circuits, modules, etc., to more evenly distribute heat generation across these various components and/or to increase the amount of processing that can be performed given a thermal or power envelope."

As indicated above, even the Examiner does not specifically argue that Orenstien says or even remotely suggests anything at all about any kind of "throughput metric", much less moving a plurality of executing computer processing jobs amongst a plurality of computer processor cores to improve a throughput metric, as presently recited. Rather, the Examiner merely argues that Orenstien discloses the use of a "performance metric" (which term apparently is not used in Orenstien itself). Moreover, the specific portions of Orenstien upon which the Examiner relies clearly indicate that the "performance metric" referenced by the Examiner pertains to temperature or power consumption, not throughput.

In the last sentence of the Examiner's argument (quoted above), the Examiner asserts, "The result of this migration based on the monitoring of the performance measurement is an increase in throughput or clock frequency for each core." In view of the foregoing remarks, it can be readily appreciated that this statement is not exactly

correct. In some of its embodiments, Orenstien apparently adjusts clock frequency for individual processors in an attempt to achieve certain thermal or power-consumption targets. It is assumed that when the clock frequency is increased for a particular processor, the throughput of that processor correspondingly increases. However, as noted above, Orenstien does not generate, use or evaluate any "throughput metric", and clearly does not move a plurality of executing computer processing jobs amongst a plurality of computer processor cores to improve any throughput metric, as presently recited.

Because this feature is not suggested by any permissible combination of claim 1 of the '147 patent and Orenstien, the present double-patenting rejection is believed to be improper.

Double-Patenting Rejection over claim 14 of the '147 patent in view of Orenstien

Claim 7

Independent claim 7 is directed to a method for operating multiple processor cores, in which multiple computer processor cores are placed on a single semiconductor die, in which at least two computer processor cores differ in processing performance, and in which all execute the same instruction set. Performance is measured for each of a set of computer processing jobs hosted amongst the plurality of computer processor cores, and individual ones of the computer processing jobs are transferred amongst targeted ones of the computer processor cores to improve a throughput metric.

The foregoing combination of features is not disclosed or suggested by the asserted combination. For instance, no permissible combination of claim 14 of the '147 patent and Orenstien would have disclosed or suggested at least the feature of transferring individual ones of a plurality of computer processing jobs amongst targeted ones of a plurality of computer processor cores to improve a throughput metric.

The Examiner acknowledges that this feature is neither disclosed nor suggested by claim 14 of the '147 patent, but then goes on to argue that Orenstien discloses such a feature. Specifically, the Examiner argues that,

"Orenstien teaches a computer system monitoring the operational activity of its dual-core processor 200 to provide load balancing between the cores by migrating processes from one core to the other based on a performance metric ([0017], [0022]-[0023]). Orenstien teaches that the performance metric, for example, can be based on power, or load for load balancing, or other things (page 2, last line of [0021]). The result of this migration based on the monitoring of the performance measurement is an increase in throughput or clock frequency for each core ([0020], [0026], [0037], [0039])." [Emphasis in original]

In this regard, Orenstien discusses techniques for modifying certain parameters pertaining to process execution based on power-consumption and thermal considerations. However, it does not say or even suggest anything at all about the above-referenced feature of the present invention.

At the outset, it is noted that even the Examiner does not identify a single "throughput metric" that is even allegedly disclosed or suggested in Orenstien, and Appellants are unable to find any such disclosure in Orenstien. Accordingly, no permissible combination of claim 14 in the '147 patent and Orenstien possibly could have disclosed or suggested the presently recited feature of transferring individual ones of a plurality of computer processing jobs amongst targeted ones of a plurality of computer processor cores to improve a throughput metric.

In addition, all of the specific portions of Orenstien cited by the Examiner pertain to the modification of processing parameters (e.g., individual processor power, individual processor clock frequency and/or processing load distribution) based on *power consumption* and/or *temperature*. None of these specifically cited portions of Orenstien says or even remotely suggests anything at all about any “throughput metric”, as presently recited.

The relevant parts of those specifically cited portions of Orenstien are briefly summarized as follows: paragraph [0017] discusses redistributing processing tasks based on power-consumption and/or thermal considerations; paragraphs [0022]-[0023] generally discuss different kinds of processors and processor cores; the last line of paragraph [0021] talks about using a “power consumption metric” in order to level the load between processors; paragraph [0020] discusses the use of a power-consumption or temperature metric for load-balancing purposes; paragraph [0026] mentions the use of a “power monitor 260”; paragraph [0037] discusses the testing of certain conditions pertaining to power consumption and then controlling voltage and/or frequency to the processors based on the results; and paragraph [0039] discusses:

“[shifting processing] between units, cores, circuits, modules, etc., to more evenly distribute heat generation across these various components and/or to increase the amount of processing that can be performed given a thermal or power envelope.”

As indicated above, even the Examiner does not specifically argue that Orenstien says anything at all about any kind of “throughput metric”, much less transferring individual ones of a plurality of computer processing jobs amongst targeted ones of a plurality of computer processor cores to improve a throughput metric, as presently recited. Rather, the Examiner merely argues that Orenstien discloses the use of a

“performance metric” (which term apparently is not used in Orenstien itself). Moreover, the specific portions of Orenstien upon which the Examiner relies clearly indicate that the “performance metric” referenced by the Examiner pertains to temperature or power consumption, not throughput.

In the last sentence of the Examiner’s argument (quoted above), the Examiner asserts, “The result of this migration based on the monitoring of the performance measurement is an increase in throughput or clock frequency for each core.” In view of the foregoing remarks, it can be readily appreciated that this statement is not exactly correct. In some of its embodiments, Orenstien apparently adjusts clock frequency for individual processors in order to achieve thermal or power-consumption targets. It is assumed that when the clock frequency is increased for a particular processor, the throughput of that processor correspondingly increases. However, as noted above, Orenstien does not generate, use or evaluate any “throughput metric”, and clearly does not transfer individual ones of a plurality of computer processing jobs amongst targeted ones of a plurality of computer processor cores to improve a throughput metric, as presently recited.

Because this feature is not suggested by any permissible combination of claim 14 of the ‘147 patent and Orenstien, the present double-patenting rejection is believed to be improper.

Rejection under § 102(e) over Orenstien

Claim 1

Independent claim 1 is directed to a computer system that includes multiple computer processor cores in which at least two differ in processing performance, and in

which all execute the same instruction set. The computer system also includes a performance measurement and transfer mechanism that moves a plurality of executing computer processing jobs amongst the computer processor cores to improve a throughput metric.

The foregoing combination of features is not disclosed by the applied art. For instance, Orenstien does not disclose at least the feature of moving a plurality of executing computer processing jobs amongst a plurality of computer processor cores to improve a throughput metric.

In this regard, Orenstien discusses techniques for modifying certain parameters pertaining to process execution based on power-consumption and thermal considerations. However, it does not say anything at all about the above-referenced feature of the present invention.

The Examiner cites the last line of paragraph [0021], together with paragraphs [0020], [0026], [0037] and [0039] in Orenstien as allegedly showing this feature of the invention. Specifically, the Examiner argues:

“In summary of the above citations, Orenstien teaches a computer system monitoring the operational activity of its dual-core processor 200 to provide load balancing between the cores by migrating processes from one core to the other based on a performance metric. Orenstien teaches that the performance metric, for example, can be based on power, or load for load balancing, or other things. The result of this migration based on the monitoring of the performance measurement is an increase in throughput or clock frequency for each core.” [Emphasis in original]

At the outset, it is noted that even the Examiner does not identify a single “throughput metric” that is even allegedly disclosed in Orenstien, and Appellants are unable to find any disclosure of any such metric in Orenstien. Accordingly, Orenstien could not possibly have disclosed the presently recited feature of moving a plurality of

executing computer processing jobs amongst a plurality of computer processor cores to improve a throughput metric.

In addition, all of the specific portions of Orenstien cited by the Examiner pertain to the modification of processing parameters (e.g., individual processor power, individual processor clock frequency and/or processing load distribution) based on *power consumption* and/or *temperature*. None of these specifically cited portions of Orenstien says anything at all about any "throughput metric", as presently recited.

The relevant parts of those specifically cited portions of Orenstien are briefly summarized as follows: the last line of paragraph [0021] talks about using a "power consumption metric" in order to level the load between processors; paragraph [0020] discusses the use of a power-consumption or temperature metric for load-balancing purposes; paragraph [0026] mentions the use of a "power monitor 260"; paragraph [0037] discusses the testing of certain conditions pertaining to power consumption and then controlling voltage and/or frequency to the processors based on the results; and paragraph [0039] discusses:

"[shifting processing] between units, cores, circuits, modules, etc., to more evenly distribute heat generation across these various components and/or to increase the amount of processing that can be performed given a thermal or power envelope."

As indicated above, even the Examiner does not specifically argue that Orenstien says anything at all about any kind of "throughput metric", much less moving a plurality of executing computer processing jobs amongst a plurality of computer processor cores to improve a throughput metric, as presently recited. Rather, the Examiner merely argues that Orenstien discloses the use of a "performance metric" (which term apparently is not used in Orenstien itself). Moreover, the specific portions of Orenstien

upon which the Examiner relies clearly indicate that the “performance metric” referenced by the Examiner pertains to temperature or power consumption, not throughput.

In the last sentence of the Examiner’s argument (quoted above), the Examiner asserts, “The result of this migration based on the monitoring of the performance measurement is an increase in throughput or clock frequency for each core.” In view of the foregoing remarks, it can be readily appreciated that this statement is not exactly correct. In some embodiments, Orenstien apparently adjusts clock frequency for individual processors in order to achieve thermal or power-consumption targets. It is assumed that when the clock frequency is increased for a particular processor, the throughput of that processor correspondingly increases. However, as noted above, Orenstien does not use any throughput metric, and clearly does not move a plurality of executing computer processing jobs amongst a plurality of computer processor cores to improve any throughput metric, as presently recited.

Lacking the above-referenced feature of the invention, Orenstien could not possibly have anticipated independent claim 1.

Accordingly, independent claim 1, together with its dependent claims, is believed to be allowable over the applied art.

Claim 7

Independent claim 7 is directed to a method for operating multiple processor cores, in which multiple computer processor cores are placed on a single semiconductor die, in which at least two computer processor cores differ in processing performance, and in which all execute the same instruction set. Performance is

measured for each of a set of computer processing jobs hosted amongst the plurality of computer processor cores, and individual ones of the computer processing jobs are transferred amongst targeted ones of the computer processor cores to improve a throughput metric.

The foregoing combination of features is not disclosed by the applied art. For instance, Orenstien does not disclose at least the features of (1) measuring performance of each of a plurality of computer processing jobs hosted amongst the plurality of computer processor cores; or (2) transferring individual ones of a plurality of computer processing jobs amongst targeted ones of the plurality of computer processor cores to improve a throughput metric.

In this regard, Orenstien discusses techniques for modifying certain parameters pertaining to process execution based on power-consumption and thermal considerations. However, it does not say anything at all about the above-referenced features of the present invention.

With respect to the first feature noted above, the Examiner simply cites paragraph [0019] of Orenstien, without comment. However, that portion of Orenstien has been studied in detail and is only seen to discuss the monitoring of power-consumption and/or temperature information or, alternatively, some factor indicating level of activity for a particular processing unit. Neither that portion of Orenstien, nor anything else in Orenstien for that matter, says anything at all about measuring performance of individual computer processing jobs, as is presently recited.

As to the second feature noted above, the Examiner cites paragraphs [0020], [0021] and [0027] of Orenstien, noting:

"In summary of the above citations, Orenstien teaches a computer system monitoring the operational activity of its dual-core processor 200 to provide load balancing between the cores by migrating processes from one core to the other based on a performance metric. Orenstien teaches that the performance metric, for example, can be based on power, or load for load balancing, or other things. The result of this migration based on the monitoring of the performance measurement is an increase in throughput or clock frequency for each core." [Emphasis in original]

At the outset, it is noted that even the Examiner does not identify a single "throughput metric" that is even allegedly disclosed in Orenstien, and Appellants are unable to find any disclosure of any such metric in Orenstien. Accordingly, Orenstien could not possibly have disclosed the presently recited feature of transferring individual ones of a plurality of computer processing jobs amongst targeted ones of the plurality of computer processor cores to improve a throughput metric.

In addition, all of the specific portions of Orenstien cited by the Examiner pertain to the modification of processing parameters (e.g., individual processor power, individual processor clock frequency and/or processing load distribution) based on power consumption and/or temperature. None of these specifically cited portions of Orenstien says anything at all about any "throughput metric", as presently recited.

The relevant parts of those specifically cited portions of Orenstien are briefly summarized as follows: paragraph [0020] discusses the use of a power-consumption or temperature metric for load-balancing purposes; paragraph [0021] talks about swapping or rotating processes between processing units based on a "thermal or power consumption metric" and/or in order to level the load between processors; and paragraph [0027] discusses monitoring power consumption for different processors and then performing "an exchange operation" based on that information.

As indicated above, even the Examiner does not specifically argue that Orenstien says anything at all about any kind of “throughput metric”, much less transferring individual ones of a plurality of computer processing jobs amongst targeted ones of the plurality of computer processor cores to improve a throughput metric, as presently recited. Rather, the Examiner merely argues that Orenstien discloses the use of a “performance metric” (which term apparently is not used in Orenstien itself). Moreover, the specific portions of Orenstien upon which the Examiner relies clearly indicate that the “performance metric” referenced by the Examiner pertains to temperature or power consumption, not throughput.

In the last sentence of the Examiner’s argument (quoted above), the Examiner asserts, “The result of this migration based on the monitoring of the performance measurement is an increase in throughput or clock frequency for each core.” In view of the foregoing remarks, it can be readily appreciated that this statement is not exactly correct. In some embodiments, Orenstien apparently adjusts clock frequency for individual processors in order to achieve thermal or power-consumption targets. It is assumed that when the clock frequency is increased for a particular processor, the throughput of that processor correspondingly increases. However, as noted above, Orenstien does not use any throughput metric, and clearly does not transfer individual ones of a plurality of computer processing jobs amongst targeted ones of a plurality of computer processor cores to improve a throughput metric, as presently recited.

Lacking the above-referenced feature of the invention, Orenstien could not possibly have anticipated independent claim 7.

Accordingly, independent claim 7, together with its dependent claims, is believed to be allowable over the applied art.

Claims 2 and 8

Claim 2 depends from independent claim 1 and claim 8 depends from independent claim 7 (discussed above). Each recites the further limitation a periodic test to determine relative performance of different jobs on different ones of the computer processor cores. This additional feature of the invention is not disclosed by the applied art.

With respect to this feature, the Examiner simply cites paragraphs [0021] and [0027] of Orenstien, arguing, “relative performance of each core is periodically monitored and evaluated such that the load can be leveled or balanced.” However, those portions of Orenstien have been studied in detail and to not appear to say anything at all about the above-referenced feature of the invention.

Rather, paragraph [0021] of Orenstien merely talks about swapping or rotating processes between processing units based on a “thermal or power consumption metric” and/or in order to level the load between processors; and paragraph [0027] of Orenstien merely discusses monitoring power consumption for different processors and then performing “an exchange operation” based on that information.

In light of what these portions of Orenstien actually discuss, it is not believed that the above-quoted statement by the Examiner is correct. Such portions of Orenstien do not say anything at all about monitoring or evaluating relative performance for load leveling or load balancing; instead, any load leveling or load balancing in Orenstien

appears to be based solely on thermal or power-consumption characteristics, not performance.

For these additional reasons, claims 2 and 8 are believed to be allowable over the applied art.

Claim 19

Claim 19 depends from independent claim 1 (discussed above) and recites the further limitation that the performance measurement and transfer mechanism periodically transfers the executing computer processing jobs to a new assignment amongst the plurality of computer processor cores, collects performance statistics about execution at the new assignment, and then determines whether to reassign the executing computer processing jobs to different computer processor cores based on the performance statistics collected. This additional feature of the invention is not disclosed by the applied art.

With respect to this feature, the Examiner simply cites paragraphs [0021] and [0027] of Orenstien, arguing, "relative performance of each core is periodically monitored and evaluated such that the load can be leveled or balanced via process migration between cores." However, those portions of Orenstien have been studied in detail and to not appear to say anything at all about the above-referenced feature of the invention.

Rather, paragraph [0021] of Orenstien merely talks about swapping or rotating processes between processing units based on a "thermal or power consumption metric" and/or in order to level the load between processors; and paragraph [0027] of Orenstien

merely discusses monitoring power consumption for different processors and then performing “an exchange operation” based on that information.

In light of what these portions of Orenstien actually discuss, it is not believed that the above-quoted statement by the Examiner is correct. Such portions of Orenstien do not say anything at all about monitoring or evaluating relative performance for load leveling or load balancing; instead, any load leveling or load balancing in Orenstien appears to be based solely on thermal or power-consumption characteristics, not performance.

In any event, neither of these specifically cited portions of Orenstien, nor anything else in Orenstien for that matter, discloses the presently recited feature of periodically transferring executing computer processing jobs to a new assignment amongst the plurality of computer processor cores, collecting performance statistics about execution at the new assignment, and then determining whether to reassign the executing computer processing jobs to different computer processor cores based on the performance statistics collected. In fact, even the Examiner has not specifically argued that Orenstien discloses this feature of the invention.

For these additional reasons, claim 19, together with its dependent claims, is believed to be allowable over the applied art.

Rejection under § 103(a) over Orenstien in view of Schmeck

Claim 18

Claim 18 depends from independent claim 1 (discussed above) and recites the further limitation that the performance measurement and transfer mechanism (which, among other things, moves a plurality of executing computer processing jobs amongst a

plurality of computer processor cores to improve a throughput metric) maximizes total system throughput. This additional feature of the invention is not disclosed or suggested by the applied art.

The Examiner acknowledges that Orenstien does not disclose or suggest this feature of the invention, but argues:

"However, Schmeck teaches that a single processor containing multiple clusters can be optimized and that the goal is to maximize the total number of instructions per second (throughput) (see page 157 under Section 3.5: Optimization): It would have been obvious to one of ordinary skill in the art at the time the invention was made to modify Orenstien such that the total system throughput is maximized because it would provide the predicted result of improved performance of the system."

At the outset, Appellants note that the portion of Schmeck cited by the Examiner appears to be discussing ASIC (application specific integrated circuit) hardware design, rather than saying anything at all about moving computer processing jobs in real time. Accordingly, Schmeck fails to say anything at all about a performance measurement and transfer mechanism, as presently recited, much less about the specific feature of such a mechanism recited in claim 18. Furthermore, because Schmeck is directed to a significantly different problem (ASIC design) than is addressed by Orenstien (real-time thermal management), there would have been absolutely no reason that one of ordinary skill in the art would have been motivated to combine these two very different references in any manner whatsoever, much less as currently argued by the Examiner.

For these additional reasons, claim 18 is believed to be allowable over the applied art.

Claim 23

Claim 23 depends from independent claim 1 (discussed above) and recites the further limitation that the throughput metric (which is improved by moving a plurality of executing computer processing jobs amongst the plurality of computer processor cores) comprises a number of instructions per second. This additional feature of the invention is not disclosed or suggested by the applied art.

The Examiner acknowledges that Orenstien does not disclose or suggest this feature of the invention, but argues:

"However, Orenstien does teach improving a throughput metric such as clock frequency of the processor. In addition, Schmeck teaches a single processor containing multiple clusters that can be optimized by maximizing the total number of instructions per second (see page 157 under Section 3.5: Optimization). It would have been obvious to one of ordinary skill in the art at the time the invention was made to modify Orenstien such that the throughput metric would comprise of a number of instructions per second. The suggestion/motivation for doing so would have been to provide the predicted result of having a metric that can be used to optimize and maximize the system (see page 157 under Section 3.5: Optimization)."

In response, Appellants first note that the Examiner has cited nothing to indicate that Orenstien teaches, "improving a throughput metric such as clock frequency of the processor." Rather, as discussed above in connection with independent claim 1, Orenstien merely teaches (in some of its embodiments) adjusting clock frequency in an attempt to modify power-consumption or temperature values.

Second, the portion of Schmeck cited by the Examiner appears to be discussing ASIC (application specific integrated circuit) hardware design, rather than saying anything at all about moving computer processing jobs in real time to improve a throughput metric. Accordingly, Schmeck fails to say anything at all about such a throughput metric comprising a number of instructions per second, as presently recited.

Third, Schmeck is directed to a significantly different problem (ASIC design) than is addressed by Orenstien (real-time thermal management). As a result, there would have been absolutely no reason that one of ordinary skill in the art would have been motivated to combine these two very different references in any manner whatsoever, much less as currently argued by the Examiner.

For these additional reasons, claim 23 is believed to be allowable over the applied art.

Rejection under § 103(a) over Orenstien in view of Matoba

Claim 20

Claim 20 depends from claim 19 (discussed above) and recites the further limitation that the performance measurement and transfer mechanism swaps execution of the executing computer processing jobs between the computer processor cores for a period of time, monitoring resulting performance, and then builds a table with relative performances of jobs on different types of cores. This additional feature of the invention is not disclosed or suggested by the applied art.

The Examiner acknowledges that Orenstien does not disclose or suggest this feature of the invention, but argues:

"However, Matoba teaches switching CPUs based on load state of each CPU and having a process management table 23 that collects and manages the information regarding the allocation (col. 7, lines 56-67, col. 8, lines 31-42, col. 9, lines 4-19). Orenstien and Matoba are analogous art because they both are related to power saving in a parallel processing environment. It would have been obvious to one of ordinary skill in the art at the time the invention was made to modify Orenstien such that it would include a table to manage and organize the information related to the monitoring functions and performance results. The suggestion/motivation for doing so would have been to provide the predicted result of improving the management and organization of data."

In response, Appellants note that Matoba concerns a computer having multiple CPUs and power-saving functionality, e.g., for use when the computer switches to battery power, that controls which CPUs are operable. See, e.g., Matoba's Title and Abstract.

The specific portions of Matoba cited by the Examiner merely talk about transferring a process to a second CPU when a first CPU is to be halted and resuming operation of the CPU. They say nothing at all about swapping execution of any executing computer processing jobs between computer processor cores for a period of time, monitoring resulting performance, and then building a table with relative performances of jobs on different types of cores, as presently recited.

In fact, they do not even indicate that CPUs are switched, "based on load state of each CPU," as asserted by the Examiner. Rather, the halting and resuming of a CPU in Matoba appears to be based on the current power mode.

Matoba's process management table 23, referenced by the Examiner, is merely described as storing the state of a running process, so that it can be transferred in an effective manner to a different CPU, when necessary.

In short, no permissible combination of Orenstien and Matoba would have disclosed or suggested the above-reference feature of the invention.

For these additional reasons, claim 20, together with its dependent claim 21, is believed to be allowable over the applied art.

Claim 21

Claim 21 depends from claim 20 (discussed above) and recites the further limitation that the jobs are reassigned based on the relative performances, by assigning

jobs that benefited most from large complex processor cores to said large complex processor cores. This additional feature of the invention is not disclosed or suggested by the applied art.

As to this feature, the Examiner argues:

"Orenstien teaches wherein the jobs are reassigned based on the relative performances, by assigning jobs that benefited most from large complex processor cores to said large complex processor cores (the monitor value may be a simple or complex activity factor that reflects the operational activity of a particular processing unit; migration is done based on monitor value) ([0019])"

In response, Appellants note that migration based on a general "activity factor that reflects the operational activity of a particular processing unit" is significantly different than reassigning jobs based on relative performances of jobs on different types of cores, by assigning jobs that benefited most from large complex processor cores to said large complex processor cores, as presently recited. In fact, even the Examiner does not specifically argue that this particular feature of the invention is disclosed or suggested by Orenstien or Matoba.

For these additional reasons, claim 21 is believed to be allowable over the applied art.

VIII. CONCLUDING REMARKS

As Appellants have shown above, for a number of reasons, nothing in the cited references discloses, teaches, or suggests the invention recited by the claims on appeal. Appellants therefore respectfully submit that the claimed invention is patentably distinct over the applied art.

In view of the foregoing remarks, Appellants respectfully request that the rejection of claims and, 5-8, 11, 12, 15 and 17-24 be reversed and a Notice of Allowance issued.

Respectfully submitted,

JOSEPH G. SWAN, A PROFESSIONAL CORP.

Dated: June 4, 2008

By /Joseph G. Swan/
Joseph G. Swan
Registration No. 41,338

CLAIMS APPENDIX

Claims on Appeal

1. A computer system, comprising:
a plurality of computer processor cores in which at least two differ in processing performance, and in which all execute the same instruction set; and
a performance measurement and transfer mechanism that moves a plurality of executing computer processing jobs amongst the plurality of computer processor cores to improve a throughput metric.
2. The computer system of claim 1, further comprising:
at least one of an operating system hosted on the plurality of computer processor cores, firmware, and special-purpose hardware that includes the performance measurement and transfer mechanism, and that provides for a periodic test to determine relative performance of different jobs on different ones of the computer processor cores.
3. (Canceled)
4. (Canceled)
5. The computer system of claim 1, further comprising:

at least one of an operating system hosted on the plurality of computer processor cores, firmware, and special-purpose hardware that includes the performance measurement and transfer mechanism, and that provides for a periodic test of operating states within each of the computer processor cores in making a decision as to where to place a given processing software workload, wherein said operating states are dependent on at least one of the operating voltage and clock frequency of a corresponding one of the plurality of computer processor cores.

6. The computer system of claim 1, further comprising:

at least one of an operating system hosted on the plurality of computer processor cores, firmware, and special-purpose hardware that includes the performance measurement and transfer mechanism, and that provides for a periodic test of operating states within each of the computer processor cores in making a decision as to where to place a given processing software workload, wherein said operating states are dependent on run-time re-configuration of hardware structures of corresponding ones of the plurality of computer processor cores.

7. A method for operating multiple processor cores, comprising:

placing a plurality of computer processor cores on a single semiconductor die, in which at least two computer processor cores differ in processing performance, and in which all execute the same instruction set;

measuring performance of each of a plurality of computer processing jobs hosted amongst the plurality of computer processor cores; and

transferring individual ones of said plurality of computer processing jobs amongst targeted ones of said plurality of computer processor cores to improve a throughput metric.

8. The method of claim 7, further comprising:

providing for a periodic test to determine relative performance of different jobs on different ones of the computer processor cores.

9. (Canceled)

10. (Canceled)

11. The method of claim 7, further comprising:

providing for a periodic test of operating states within each of the computer processor cores in making a decision as to where to place a given processing software workload, wherein said operating states are dependent on at least one of the operating voltage and clock frequency of a corresponding one of the plurality of computer processor cores.

12. The method of claim 7, further comprising:

providing for a periodic test of operating states within each of the computer processor cores in making a decision as to where to place a given processing software

workload, wherein said operating states are dependent on run-time re-configuration of hardware structures of corresponding ones of the plurality of computer processor cores.

13. (Canceled)

14. (Canceled)

15. The method of claim 7, further comprising:
associating workloads for execution on specific processor cores based on at least one of user and application hints.

16. (Canceled)

17. The computer system of claim 1, further comprising at least one of an operating system hosted on the plurality of computer processor cores, firmware, and special-purpose hardware that includes the performance measurement and transfer mechanism.

18. The computer system of claim 1, wherein the performance measurement and transfer mechanism maximizes total system throughput.

19. The computer system of claim 1, wherein the performance measurement and transfer mechanism periodically transfers the executing computer processing jobs

to a new assignment amongst the plurality of computer processor cores, collects performance statistics about execution at the new assignment, and then determines whether to reassign the executing computer processing jobs to different computer processor cores based on the performance statistics collected.

20. The computer system of claim 19, wherein the performance measurement and transfer mechanism swaps execution of the executing computer processing jobs between the computer processor cores for a period of time, monitoring resulting performance, and then builds a table with relative performances of jobs on different types of cores.

21. The computer system of claim 20, wherein the jobs are reassigned based on the relative performances, by assigning jobs that benefited most from large complex processor cores to said large complex processor cores.

22. The computer system of claim 19, wherein the determination of whether to reassign the jobs to different computer processor cores also is based on at least one of user-defined or workload-defined metrics.

23. The computer system of claim 1, wherein the throughput metric comprises a number of instructions per second.

24. The computer system of claim 1, wherein movement of the executing computer processing jobs is constrained to occur only at operating system time slice intervals.

EVIDENCE APPENDIX

None

RELATED PROCEEDINGS APPENDIX

None